

# Improved model-based clustering performance using Bayesian initialization averaging.

Adrian O'Hagan · Arthur White

Received: June 26th 2016 / Accepted: TBD

**Abstract** The Expectation-Maximization (EM) algorithm is a commonly used method for finding the maximum likelihood estimates of the parameters in a mixture model via coordinate ascent. A serious pitfall with the algorithm is that in the case of multimodal likelihood functions, it can get trapped at a local maximum. This problem often occurs when sub-optimal starting values are used to initialize the algorithm. Bayesian initialization averaging (BIA) is proposed as an ensemble method to generate high quality starting values for the EM algorithm. Competing sets of trial starting values are combined as a weighted average, which is then used as the starting position for a full EM run. The method can also be extended to variational Bayes (VB) methods, a class of algorithm similar to EM that is based on an approximation of the model posterior. The BIA method is demonstrated on real continuous, categorical and network data sets, and the convergent log-likelihoods and associated clustering solutions presented. These compare favorably with the output produced using competing initialization methods such as random starts, hierarchical clustering and deterministic annealing, with the highest available maximum likelihood estimates obtained in a higher percentage of cases, at reasonable computational cost. The implications of the different clustering solutions obtained by local maxima are also discussed.

**Keywords** Bayesian model averaging · Expectation-maximization algorithm · Finite mixture models · Hierarchical clustering · Model-based clustering · Multimodal likelihood.

## 1 Introduction

The Expectation-Maximization (EM) algorithm is a commonly used coordinate ascent method for finding the maximum likelihood estimates for the parameters in a mixture model (Dempster et al,

---

A. O'Hagan  
School of Mathematics and Statistics and The Insight Centre for Data Analytics, University College Dublin, Ireland  
Tel.: +353-1716-2428  
Fax: N/A  
E-mail: adrian.ohagan@ucd.ie

A. White  
School of Computer Science and Statistics, Lloyd Institute, Trinity College Dublin, Ireland

1977). To initialize the algorithm, starting values are chosen for the allocations of observations to the available number of clusters and these lead to the initial calculation of the model parameters. Often this starting allocation is done on a random basis. However, once these parameter values have been specified, the algorithm proceeds in an entirely deterministic fashion.

An advantage of the EM algorithm is that it does not produce decreases in the likelihood function. A drawback of the algorithm is that it can converge to or get trapped for many iterations at a local maximum, leading to failure to reach the global maximum and resulting in an inferior clustering solution. Zhou and Lange (2010) describe this multimodality of the likelihood function as "one of the curses of statistics". Various convergence criteria can be used to stop the algorithm. However these criteria tend to characterize a lack of progress of the algorithm, rather than the fact that the global maximum has been reached.


Many adaptations to the EM algorithm have been proposed that address this issue by altering its form. These include, among others, an EM algorithm featuring maximization using the Newton-Raphson method (Redner and Walker, 1984); the Classification EM (CEM) algorithm (Biernacki et al, 2003; Celeux and Govaert, 1992); the Moment Matching EM algorithm (Karlis and Xekalaki, 2003); the Annealing EM algorithm (Zhou and Lange, 2010); a hybrid EM/Gauss-Newton algorithm (Aitkin and Aitkin, 1996); the Expectation Conditional Maximization (ECM) algorithm (Meng and Rubin, 1992); the emEM algorithm (Biernacki et al, 2003); the Multicycle EM algorithm (Meng and Rubin, 1993) and the Sparse EM algorithm (Neal and Hinton, 1999).

Alternatively this problem can be addressed by identifying high quality, as opposed to random, starting values for the algorithm in an attempt to optimize the eventual clustering solution. Recent developments of this approach include O'Hagan et al (2012), who develop a pyramid burn-in scheme, and Baudry and Celeux (2015), who propose a recursive initialisation strategy based on splitting a selected component from a previous clustering solution, in particular with a view to performing robust model selection.

Bayesian initialisation averaging (BIA) is proposed as a new and complementary method in this spirit. To generate starting values using BIA, a trial number of E-steps and M-steps are run on each of a sequence of random starting positions to give an updated sequence of potential starting positions. Using the likelihood associated with each updated starting position, a corresponding weight is calculated and a new single overall weighted average starting position is formed. The EM algorithm is then run to convergence from this juncture.

For some types of model, the log-likelihood is intractable, and inference using EM is not possible. We demonstrate our approach, in a Bayesian setting, for one such model using variational Bayes (VB) methods, an EM-type algorithm based on an approximation of the model posterior. Related algorithms for this type of scenario include iterated conditional models (Besag, 1986), the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (Byrd et al, 1995; Zhu et al, 1997) and sequential Monte Carlo for VB (McGrory and Ahfock, 2014). The latter shares some features in common with annealing and pyramid burn-in.

As an illustrative example of BIA, consider Figure 1, which is based on a clustering model using a VB algorithm. (See Sections 2.3 and 3.2.1 for specific details). The dashed lines show the value of the (lower bound to the) log-posterior at each iteration using two sets of randomly generated suboptimal starting positions. Using these starting positions, the algorithm converges to distinct solutions which are similarly valued in terms of log-posterior. Using BIA, it is possible to obtain a new, superior set of starting values, based on a weighted combination of the value of the two positions after running the algorithm for just 20 iterations. Hence a superior clustering solution in terms of log-posterior (the solid line) may be obtained.



**Fig. 1** Illustrative example of the Bayesian initialisation averaging approach applied to a clustering model for network data using an EM-like algorithm. The dashed lines show the value of the lower bound to the log-posterior at each iteration using suboptimal starting positions; by using BIA, a superior clustering solution (the solid line) is obtained.

BIA starting values are used as an initialization method to cluster a variety of continuous, categorical and network data examples using mixtures of Gaussians, latent class analysis (LCA) and the stochastic blockmodel (SBM), implemented via the EM (mixture of Gaussians and LCA) and VB (SBM). Even for the more complicated SBM approach where the likelihood can only be approximated, the proposed BIA method fares well. We demonstrate the approach on real and simulated data sets in comparison to competing methodologies such as random starts, hierarchical clustering, pyramid burn-in schemes and deterministic annealing.

Section 2 presents the data sets used as motivating examples to illustrate the benefits of the proposed method. Section 3 gives an overview of how the EM algorithm is used for fitting mixture distributions. The calculations for the E-step and M-step are provided for different types of appli-

cation. A description of how to apply BIA to generate EM algorithm starting values is detailed. A label switching problem that is encountered, as well as a means of overcoming it, are documented. In Section 4 the outcomes of the methods tested on the data sets referenced in Section 2 are described. Convergent log-likelihood plots and clustering solutions are presented for BIA and its competing methodologies in each case. Section 5 details the conclusions reached and summarizes the main findings.

## 2 Illustrative Data Sets

Three types of data set are used as motivating examples in testing the BIA method. The first type is multivariate continuous data (the *AIS* data set and a simulated data set taken from Baudry et al (2015)); the second type is multivariate categorical data (the *Carcinoma* and *Alzheimer’s* data sets); the third type is network data (the *Karate* data set). These are described below.

### 2.1 Australian Institute of Sport (AIS) data and simulated data

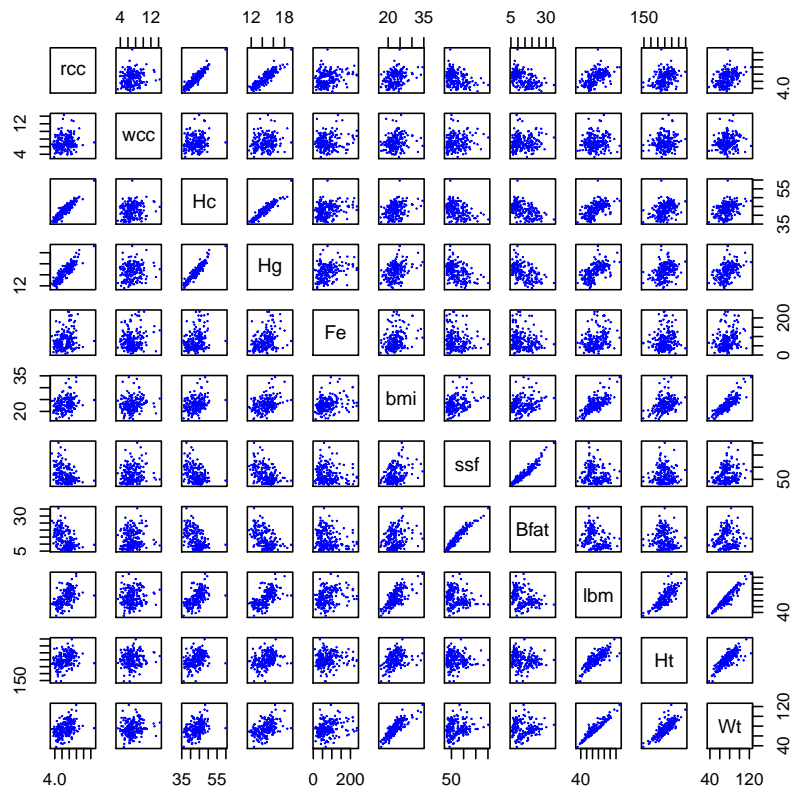
The *Australian Institute of Sport* data set contains biometric observations on  $n = 202$  Australian athletes (102 males and 100 females) at the Australian Institute of Sport (Cook and Weisberg, 1994). The full data set contains 13 variables (11 continuous and 2 discrete) and all continuous variables are included in the analysis: red cell count (rcc), white cell count (wcc), hematocrit (Hc), hemoglobin (Hg), plasma ferritin (Fe), body mass index (bmi), sum of skin folds (ssf), body fat percentage (Bfat), lean body mass (lbm), height (Ht) and weight (Wt) (Cook and Weisberg, 1994) (see Figure 2).

A further simulated data set taken from Baudry et al (2015) experiment 3 is tested, namely 200 observations simulated from a mixture of three multivariate Gaussian components with diagonal covariance structure. The second and third horizontal components (plotted as red crosses) overlap, giving the appearance of forming a single cluster, as seen in Figure 3. The first component, plotted as black circles, is vertical in nature. For this data set, both a 2 group and a 3 group solution are relevant. 100 such data sets are randomly simulated for the purposes of the study.

### 2.2 Carcinoma and Alzheimer’s data

Both the *Carcinoma* and *Alzheimer’s* data sets consist of multivariate binary data, to which LCA is applied. The *Carcinoma* data set contains dichotomous ratings by  $m = 7$  pathologists of  $n = 118$  slides for the presence or absence of carcinoma in the uterine cervix. The data set is available in the **R** package **poLCA** (Linzer and Lewis, 2011) and is described at length by Agresti (2002, Section 13.2). As per the simple percentage diagnosed summary for each doctor in Table 1, the data clearly demonstrate that some pathologists are more prone to diagnosing the presence of carcinoma than others. Using LCA we aim to identify sets of observations for which they agree and disagree. The approach of Zhou and Lange (2010) is followed and the case of  $G = 4$  classes considered.

The *Alzheimer’s* data set documents the presence or absence of  $m = 6$  symptoms in  $n = 240$  patients diagnosed with early onset Alzheimer’s disease, conducted in Saint James’ Hospital, Dublin, Ireland. This data is available in the **R** package **BayesLCA** (White and Murphy, 2014), and was previously analyzed by Moran et al (2004) and Walsh (2006), where the goal was to identify whether



**Fig. 2** Pairs plot of the *Australian Institute of Sport (AIS)* data.


**Table 1** Percentage of slides identified as having cancer present by the seven different doctors in the *Carcinoma* study.

Pathologist	A	B	C	D	E	F	G
Detected (%)	0.56	0.67	0.38	0.27	0.60	0.21	0.56

groups of patients with the condition presented distinctive patterns of symptoms. The case where  $G = 3$  latent classes are assumed to be present in the data is investigated.

### 2.3 Karate data

Zachary's Karate club data set records the social network of friendships between  $n = 34$  members of a karate club at a United States university in the 1970s (Zachary, 1977). While friendships within the network initially arose organically, a dispute between a part time instructor and the club president led to two political factions developing, with the members taking the side of either the




**Fig. 3** Pairs plot of the simulated mixture of Gaussians data.

president or the instructor, ultimately leading to the formation of a new organisation. A  $G = 4$  group SBM is applied to this data. The network is visualised in Figure 4 using the Fruchterman-Reingold (Fruchterman and Reingold, 1991) layout algorithm available in the **R** package **igraph** (Csardi and Nepusz, 2006).

### 3 Methods

#### 3.1 The EM Algorithm

The EM algorithm is a powerful computational technique for finding the maximum likelihood estimates (MLEs) of the parameters in a mixture model when there are no closed-form maximum likelihood estimates, or when the data are incomplete (McLachlan and Krishnan, 1997). The EM algorithm was introduced by Dempster et al (1977) and has many different applications, including cluster analysis, censored data modelling, mixed models and factor analysis. In model-based



**Fig. 4** Plot of the *Karate* network data set.

clustering, the EM algorithm maximizes the expected complete data log-likelihood to produce updated parameter values. Ultimately this process maximizes the observed data log-likelihood. The algorithm involves two steps, the E or expectation step and the M or maximization step. Initial starting values are required for the model parameters and the E and M steps are then iterated until convergence. In terms of initializing the algorithm the aim is to generate good-quality starting values that will lead to the algorithm converging to the highest possible log-likelihood.

### 3.2 Fitting mixture models using the EM Algorithm

Let  $\mathbf{X}$  be the matrix of all  $n$  observations and the  $m$ -dimensional vector  $\mathbf{x}_i = x_{i1}, \dots, x_{im}$  denote the value of the  $i^{th}$  observation, for  $i = 1, 2, \dots, n$ . Each observation belongs to one of  $G$  groups.  $\mathbf{z}_{ig} = 1$  when observation  $i$  belongs to group  $g$  and  $\mathbf{z}_{ig} = 0$  otherwise. Group labels are not known in advance.  $\mathbf{Z}$  is an  $(n \times G)$  classification matrix of  $\mathbf{z}_{ig}$  values for the matrix  $\mathbf{X}$ .  $\tau_g$  is the probability that  $\mathbf{x}_i$  belongs to group  $g$ .  $\theta_g$  denotes the distributional parameters for group  $g$ . The composition of the parameter set  $\theta$  varies depending on the application at hand; for example, in a continuous data setting where a mixture of Gaussians is fitted,  $\theta$  could represent the group means and covariance matrices. Further examples are given in Section 3.2.1.

If each observation  $\mathbf{x}_i$  belongs to one of  $G$  groups with probability  $\tau_g$  then the density of  $\mathbf{x}_i$ , conditioning on the distributional parameters  $\boldsymbol{\theta}$  and group membership probabilities  $\boldsymbol{\tau}$ , is given by:

$$\sum_{g=1}^G [\tau_g P(\mathbf{x}_i | \boldsymbol{\theta}_g)] \quad (1)$$

The observed likelihood is the product of the densities of the  $\mathbf{x}_i$ , conditioning on  $\boldsymbol{\theta}$  and  $\boldsymbol{\tau}$ :

$$L = P(\mathbf{X} | \boldsymbol{\tau}, \boldsymbol{\theta}) = \prod_{i=1}^n \sum_{g=1}^G [\tau_g P(\mathbf{x}_i | \boldsymbol{\theta}_g)]. \quad (2)$$

This function is difficult to work with as its natural logarithm cannot be differentiated in a straightforward manner. However, introducing the  $\mathbf{Z}$  classification matrix, the complete data likelihood  $L_c$  can be expressed as:

$$L_c = P(\mathbf{X}, \mathbf{Z} | \boldsymbol{\tau}, \boldsymbol{\theta}) = \prod_{i=1}^n \prod_{g=1}^G [\tau_g P(\mathbf{x}_i | \boldsymbol{\theta}_g)]^{z_{ig}}. \quad (3)$$

The complete data log-likelihood  $l_c$  can then be written as:

$$l_c = \sum_{i=1}^n \sum_{g=1}^G z_{ig} [\log \tau_g + \log P(\mathbf{x}_i | \boldsymbol{\theta}_g)]. \quad (4)$$

The EM algorithm maximizes the observed log-likelihood, (2), according to the following process:

- (i) **Starting Values:** Set  $\mathbf{Z}^{(0)}$  at iteration  $t = 0$ . Calculate the values of  $\boldsymbol{\tau}^{(0)}$  and  $\boldsymbol{\theta}^{(0)}$  based on  $\mathbf{Z}^{(0)}$ .
- (ii) **E-Step:** Evaluate  $Q^{(t+1)} = E[l_c(\boldsymbol{\theta}, \boldsymbol{\tau}, \mathbf{Z}) | \boldsymbol{\theta}^{(t)}, \boldsymbol{\tau}^{(t)}]$ , which means estimating  $E(z_{ig})$ :

$$\hat{e}_{ig}^{(t+1)} = E[z_{ig}^{(t+1)}] = \frac{\tau_g^{(t)} P(\mathbf{x}_i | \boldsymbol{\theta}_g^{(t)})}{\sum_{g'=1}^G \tau_{g'}^{(t)} P(\mathbf{x}_i | \boldsymbol{\theta}_{g'}^{(t)})}. \quad (5)$$

- (iii) **M-Step:** Maximize  $Q^{(t+1)}$  (Equation 6) with respect to the group membership probabilities,  $\boldsymbol{\tau}$ , and the distributional parameters of the observed data,  $\boldsymbol{\theta}$ . This produces new values  $\boldsymbol{\tau}^{(t+1)}$  (see Equation 7) and  $\boldsymbol{\theta}^{(t+1)}$ , where the composition of  $\boldsymbol{\theta}^{(t+1)}$  varies according to the type of data:

$$Q^{(t+1)} = \sum_{i=1}^n \sum_{g=1}^G \hat{e}_{ig}^{(t+1)} [\log \tau_g + \log P(\mathbf{x}_i | \boldsymbol{\theta}_g)] \quad (6)$$

$$\hat{\tau}_g^{(t+1)} = \frac{1}{n} \sum_{i=1}^n \hat{e}_{ig}^{(t+1)}. \quad (7)$$



- (iv) **Convergence Check:** The algorithm is stopped when the log-likelihood has converged. Convergence is deemed to have been reached when the relative change in the log-likelihood is “sufficiently small”:

$$\frac{l(\boldsymbol{\theta}^{(t+1)}, \boldsymbol{\tau}^{(t+1)}) - l(\boldsymbol{\theta}^{(t)}, \boldsymbol{\tau}^{(t)})}{l(\boldsymbol{\theta}^{(t+1)}, \boldsymbol{\tau}^{(t+1)})} < \epsilon, \quad (8)$$

where  $\epsilon$  is a suitably small value specified by the user. For the continuous data and SBM applications (*AIS* and *Karate* data respectively),  $\epsilon$  was chosen as  $1 \times 10^{-5}$  to keep the method consistent with alternative approaches such as hierarchical clustering initialization in **mclust** (Fraley and Raftery, 1999), which use the same convergence criterion. For the LCA examples (*Carcinoma* and *Alzheimer’s* data),  $\epsilon$  was chosen to be  $1 \times 10^{-9}$ , in keeping with previous experiments performed by Zhou and Lange (2010). The EM algorithm parameter estimates at convergence are used as a means of clustering the data set.

A common problem with the EM algorithm is that it can get trapped at a local maximum, leading to an inferior clustering solution. A computationally intensive remedy is to re-run the algorithm from many random  $\mathbf{Z}$  starting positions to convergence and then select the solution with the highest log-likelihood. Other approaches include specifying starting values by first using some type of deterministic heuristic algorithm, or employing an annealing approach. Alternatively, it is possible to optimize starting values by combining multiple random starts together. Bayesian initialization averaging (BIA) provides a means of achieving this goal by using a weighted combination of  $\mathbf{Z}$  starting positions.

While the objective is generally to find the global log-likelihood maximum, it must be noted that this mode may not be optimal in terms of how meaningful the clustering is. In models where there are no restrictions on the covariance structure the maximum log-likelihood is actually infinite. Additionally, some spurious solutions will yield convergent log-likelihoods that are “too high”, typically in cases where additional components are fitted to capture small groups of outlying observations. McLachlan and Peel (2000) document this phenomenon for mixtures of normal distributions. For the *AIS*, *Carcinoma*, *Alzheimer’s* and *Karate* data sets, spurious solutions are not an issue due to the number of clusters considered, which are in turn guided by previously published results under competing methodologies. Overall, any reference to the “global mode” should be interpreted as referring to the highest finite log-likelihood that emerges from a meaningful clustering solution in a stable subset of the parameter space.

On a related note, it is worth recognising that simply ranking models by convergent likelihood is not on its own proof of superior clustering fit and hence a superior initialization scheme. However, for the motivating data sets presented, the clustering solutions corresponding to higher convergent likelihoods at a very minimum give valuable insight as to alternative interpretations of how the data can be clustered; which can be more intuitive or insightful. This point is analyzed by McLachlan and Peel (2000), who argue that for continuous data the inter-component sum of squared distances to the cluster mean tends to improve (decrease) as convergent log-likelihood increases.

### 3.2.1 Examples of mixture models

The following mixture models were used to cluster the data sets described in Section 2.

**Mixture of Gaussians** In the continuous data setting, where  $\mathbf{x}_i \in \mathbb{R}^m$ , a mixture of Gaussians (Fraley and Raftery, 1999) can be fitted. Here  $\boldsymbol{\theta}$  represents the group means  $\boldsymbol{\mu}$  and the group variances/covariances  $\boldsymbol{\sigma}$  or  $\boldsymbol{\Sigma}$ , depending on whether the data is univariate or multivariate respectively. Several possible covariance structures are available; see Fraley and Raftery (1999) for further details.

**Latent Class Analysis** LCA (Goodman, 1974) involves clustering of categorical (specifically nominal) data, which in the simplest case is binary in nature denoting, e.g., the presence or absence of a symptom. In this setting, each  $x_{im} \in \{0, 1\}$ , and  $\boldsymbol{\theta}$  is a  $G \times M$  matrix often referred to as the item probability parameter, that is,  $\theta_{gm}$  denotes the probability, conditional on membership of group  $g$ , that  $x_{im} = 1$ , for any  $i \in 1, \dots, n$ ,

**Stochastic Blockmodel** Social network analysis (Wasserman and Faust, 1994) is the study of how links, such as friendship, between a set of actors are formed. Here,  $\mathbf{X}$  is an  $n \times n$  matrix such that

$$x_{ij} = \begin{cases} 1 & \text{if a link exists between actors } a_i \text{ and } a_j \\ 0 & \text{otherwise.} \end{cases}$$

The SBM (Holland et al, 1983; Snijders and Nowicki, 1997; Daudin et al, 2008) introduces  $G$  latent groups underlying the network. Here,  $\boldsymbol{\theta}$  is a  $G \times G$  connectivity matrix that, conditional on their group membership, represents the probability of a link being formed between two actors. Note that inference in this setting is not as straightforward as for the mixture of Gaussians and LCA cases, and the EM algorithm cannot be used directly for this model. This is because the observed data likelihood function is not available in closed form, and an approximation is required; see Daudin et al (2008) for further details. Volant et al (2012) use a VB method to perform inference for the SBM. This approach introduces a tractable lower bound  $\mathcal{L} < l$  to the observed data log-likelihood which can itself be maximised in an EM-like iterative manner. To make use of this method requires specifying prior distributions. In the example in Section 4.5, uninformative (proper) priors were chosen.

The LCA and SBM approaches make a conditional independence assumption (Hand and Yu, 2001), whereby, conditional on group membership, the data points are assumed to be independent. Mixture models also assume conditional independence - the covariance structure models dependence between variables as opposed to data points.

### 3.3 Generating Starting Values

#### 3.3.1 Random Allocation

A simple method for initializing the EM algorithm is to randomly allocate observations to the available groups. In the examples in Section 4, unless otherwise stated, the algorithm is initialized according to the following method:  $\mathbf{z}_i^{(0)} \sim \text{Multinomial}(1, 1/G, \dots, 1/G)$ , for each  $i \in 1, \dots, n$ .

Alternatively, the algorithm can be initialized by specifying random values for the parameters  $\boldsymbol{\theta}$  and  $\boldsymbol{\tau}$ . For example, in a Bayesian setting, these could be generated from the model priors. However, following experimentation, superior performance was achieved by randomly generating values for  $\mathbf{Z}^{(0)}$ .

### 3.3.2 Hierarchical Clustering Starting Values

For multivariate continuous data sets, **mclust** utilizes hierarchical clustering to generate a starting  $\mathbf{Z}$  value for the EM algorithm. Hierarchical clustering constructs a tree-like structure known as a dendrogram to show groups of observations. The final clustering is built up over a number of steps, where similar observations are joined together. Several measures can be used to calculate the distance or dissimilarity between observations, with Euclidean distance the most common. There are also many criteria of aggregation that may be used such as the single linkage criterion, the complete linkage criterion and Ward’s method (Ward, 1963). The observations with the lowest dissimilarities are grouped together first and the process then iterates on the group level until all observations belong to a single group and the dendrogram mapping the process is constructed. Once the tree is formed, the number of groups identified as optimal depends on where the tree is cut. The dendrogram is conventionally cut where there is relatively wide range of distances over which the number of clusters in the solution does not change. An advantage of hierarchical clustering is that the number of clusters is not assumed in advance. However the starting values yielded are not necessarily optimal in terms of eventually achieving the global maximum log-likelihood. BIA represents an alternative to this initialization method for generating better-quality starting values to facilitate ultimately obtaining the highest possible convergent log-likelihood.

### 3.3.3 Burn-in Methods and Parameter Targeting

O’Hagan et al (2012) detail two novel EM algorithm initialization schemes for model-based clustering of multivariate continuous data. The “burn-in” methods proposed start with a set of candidate  $\mathbf{Z}$  matrices, run each  $\mathbf{Z}$  for a number of EM iterations, rank the  $\mathbf{Z}$ s according to likelihood and retain only a certain fraction of the candidate  $\mathbf{Z}$ s for the next set of EM runs, until only one  $\mathbf{Z}$  remains. The parameter targeting method evaluates which parameter(s) in the model are most effective in driving the likelihood uphill at consistent intervals and focuses M steps on those parameters, significantly reducing the overall computational burden. BIA represents a significantly faster and statistically more robust alternative to these schemes.

### 3.3.4 Deterministic Annealing

An alternative means of identifying a global likelihood maximum is provided by deterministic annealing (Ueda, 1998; Zhou and Lange, 2010). Under this approach, a positive tuning parameter  $\nu$  is used to flatten the likelihood function, with the adapted likelihood rendered easier to explore and ultimately identify the global maximum using the EM algorithm. If implemented correctly, the algorithm should then obtain the global maximum regardless of the initial starting values.

In the case of finite mixture models, the approach is implemented by setting

$$\hat{e}_{ig}^{(t+1)} = \frac{\left[ \tau_g^{(t)} P(\mathbf{x}_i | \boldsymbol{\theta}_g^{(t)}) \right]^{\nu^{(\kappa)}}}{\sum_{g'=1}^G \left[ \tau_{g'}^{(t)} P(\mathbf{x}_i | \boldsymbol{\theta}_{g'}^{(t)}) \right]^{\nu^{(\kappa)}}},$$

where  $\nu^0$  is initially set to some small value, e.g.,  $\nu^0 = 0.05$ , then increased every  $s$  iterations towards  $\nu^\infty = 1$ . The positively valued  $r$  determines the rate of convergence to  $\nu^\infty$ , by setting  $\nu^{(\kappa+1)} = r\nu^{(\kappa)} + (1-r)\nu^\infty$ . Note that  $\nu^0, r$  and  $s$  must all be specified by the user.

### 3.4 Bayesian Initialization Averaging

In standard statistical practice, a single model is selected from a collection of competing models as the "best" model, using a metric such as BIC, and inferences are made on the basis of this "best" model being the true model (Raftery et al, 2005). Consequently the element of model uncertainty is often ignored. As a general framework, Bayesian model averaging (BMA) can provide a mechanism to account for this uncertainty (Hoeting et al, 1999). By accounting for model uncertainty, BMA minimizes prediction risk and has also been shown to improve model prediction accuracy on average by Wintle et al (2003), Volinsky et al (1997) and Murphy and Wang (2001), among others. It does this by averaging over competing models to provide broader conclusions about parameter estimates than from using a single model. Usually BMA is used after models are fitted in order to weight over the convergent results of multiple models. Conversely, its use is now proposed as an initialization method to generate starting values for the EM algorithm in an attempt to maximize the convergent log-likelihood.

The rationale for using a BMA-like approach to initialization is as follows: by averaging candidate starting values according to their relative quality of fit at a pre-convergence stage in running the EM algorithm, the overall weighted starting value is then capable of matching or exceeding the performance of any single individual candidate starting value, as it combines information on the allocation of observations to clusters from multiple evolutions of the model-fitting process. From a computational perspective, the approach allows a large number of candidate solutions to be proposed and evaluated in an efficient manner.

The following steps are followed to apply Bayesian *initialization* averaging (BIA) to generate the EM algorithm starting values for model-based clustering. In the case of the SBM, the lower bound  $\mathcal{L}$  is used as a substitute for  $l$ .

- (i) Specify a model structure or set of candidate models.
- (ii) Run a small number of preliminary E-steps and M-steps on each of a series of randomly generated  $\mathbf{Z}$  starts to give an updated sequence of  $\mathbf{Z}$  matrices, using the model structure(s) identified in (i).
- (iii) Calculate the value of an information criterion similar in nature to the Bayesian Information Criterion (Schwarz, 1978) associated with each updated  $\mathbf{Z}$  matrix.  $BIC_{Z_j}^*$  is the value of this criterion associated with the  $j^{th}$   $\mathbf{Z}$  matrix, calculated as:

$$BIC_{Z_j}^* = -2l_{Z_j} + p_j \log(n). \quad (9)$$

where  $l_{Z_j}$  is the observed log-likelihood for the  $j^{th}$   $\mathbf{Z}$  matrix,  $p_j$  is the number of parameters in the  $j^{th}$  model and  $n$  is the number of observations. The key difference between  $BIC^*$  and the standard formulation is that the criterion is not evaluated at the maximum observed likelihood, since this quantity will not yet have been identified. Despite not operating at the maximum observed likelihood, empirically it is found that the penalty term  $p_j \log(n)$  remains an effective discriminator between models of varying complexity when assigning weights to the competing starting positions.

- (iv) Re-scale the  $BIC^*$  values by subtracting the maximum value from each of the individual elements. This renders the results computationally tractable for evaluation in the remaining steps.

- (v) Calculate the weight for the  $j^{th}$   $\mathbf{Z}$  matrix,  $w_j$ , for BIA using the formula:

$$w_j = \frac{\exp(-0.5BIC_{Z_j}^*)}{\sum_{\forall k} \exp(-0.5BIC_{Z_k}^*)}. \quad (10)$$

- (vi) Form a single new  $\mathbf{Z}$  matrix,  $\mathbf{Z}^*$ , as the weighted average of the  $\mathbf{Z}_j$  matrices using the values calculated by equation 10 as the weights:

$$\mathbf{Z}^* = \sum_{\forall j} w_j \mathbf{Z}_j. \quad (11)$$

- (vii) Run the EM algorithm to convergence using the single BIA weighted starting  $\mathbf{Z}^*$  matrix formed in (vi).

### 3.5 The Label Switching Problem

A label switching problem is encountered when applying the BIA method, prior to implementing step (vi) above. The maximum likelihood label configurations from the runs of EM algorithm are invariant to permutations of the data labels. Hence, if the BIA weighting process is applied to a large number of candidate  $\mathbf{Z}$  matrices that have undergone a preliminary set of EM steps, the result is merely a single  $\mathbf{Z}$  matrix with roughly equal membership probabilities for all observations across all available groups.

The method employed to undo the label switching problem is a “soft” label switching method. This method involves using the actual  $\mathbf{Z}$  values or “soft” group membership probabilities associated with each observation (Slonim et al, 2005). Hence observations on the boundaries between several clusters are not forced to fully belong to one of the clusters, but rather are assigned membership percentages indicating their partial membership (Rokach and Maimon, 2010). The `matchClasses` function in **R** was used to switch the labels. The `matchClasses` function takes the actual values in the array of  $\mathbf{Z}$  matrices and matches them with a fixed row in the  $\mathbf{Z}$  matrix to fix the labels. Each row of the matrix is sequentially matched versus the fixed row until the overall labelling structure is repaired. Using the actual  $\mathbf{Z}$  probabilities avoids the loss of information associated with label-switching methodologies that operate on the “hard” group labels, such as that associated with the work of Carpaneto and Toth (1980) and Nobile and Fearnside (2007). However, for completeness, the latter “hard” label-switching methodology was checked for consistency in the continuous data case and deviations in outcomes versus the “soft” approach were found to be minimal. The method is also robust in the presence of empty clusters, though this issue is not encountered for any of the motivating data sets considered. While it is true that no method for correcting for label switching is guaranteed to provide an optimal solution under all circumstances, the results presented demonstrate that the “soft” approach performs well for the real and simulated data considered.

## 4 Results


The results in this section were obtained using a varying number of initial  $\mathbf{Z}$  matrices and preliminary EM runs for each data set and both label switching methods. The BIA code was first run for a range of different values of the number of initial  $\mathbf{Z}$  matrices and the number of preliminary EM runs and the optimum combination was then used to produce the results below. For some

data sets, this number of initial  $\mathbf{Z}$  matrices and preliminary EM steps may seem computationally costly. However, superior BIA performance can still be witnessed over competing methods at a much lower computational cost by reducing the number of  $\mathbf{Z}$ s and preliminary EM steps, but not as frequently. Under the method of random starts, all randomly generated  $\mathbf{Z}$  matrices are run until full convergence of the EM algorithm. The early iterations of any EM algorithm are usually the most profitable in terms of driving the likelihood function uphill. The BIA method exploits this since it focuses computational power on these early iterations prior to weighting the processed  $\mathbf{Z}$  matrices to form a single  $\mathbf{Z}$  for use in the EM algorithm.

#### 4.1 Australian Institute of Sports (AIS) data

Figure 5 displays a histogram of the convergent log-likelihoods for the 11 continuous variables of the *AIS* data set, using BIA, with both 40 initial  $\mathbf{Z}$  matrices and 50 preliminary EM runs (Figure 5(b)) and 50 initial  $\mathbf{Z}$  matrices and 100 preliminary EM runs (Figure 5(c)) respectively. The former gives a similar distribution of convergent log-likelihood values to that obtained using the previously developed pyramid burn-in methodology developed by (O’Hagan et al, 2012) in Figure 5(a). The latter gives a superior distribution of convergent log-likelihoods. The model applied has  $G = 2$  multivariate normal components and ellipsoidal covariance structure with equal shape and volume (EEV in standard **mclust** nomenclature) as selected by BIC. The purple vertical line represents the optimal log-likelihood reached by **mclust**’s hierarchical clustering initialization with  $G = 2$  groups and EEV covariance structure. The “soft” label-switching methodology is employed in this instance, but there is minimal effect on the results of alternatively using the “hard” label-switching approach. The BIA approach is extremely fast to run, requiring only 5% of the processing time needed by the alternative schemes developed by O’Hagan et al (2012) for the *AIS* data. This shows that BIA has potential to be applied to data sets with large numbers of observations and variables, or to more complex mixture models for continuous data. For thoroughness, the BIA method was also applied to the *Virginicas*, *Galaxies* and *Hidalgo* data sets used as motivating examples by O’Hagan et al (2012). As was the case with the *AIS* data, the distribution of convergent log-likelihoods was comparable to or better than that found using the competing burn-in and parameter targeting methods, but at considerably lower computational cost.

Figure 6 displays the pairs plot for body mass index (bmi) and body fat Percentage (Bfat) under the clustering solution obtained using a hierarchical initialization of the EM algorithm (convergent log-likelihood of  $-4,743.6$ ). Compared to the optimal convergent log-likelihood that is regularly identified by the BIA method ( $-4,722.4$ ), 15 data points change cluster membership, highlighted as crosses. The change in clustering solution produces a shift in the total within-cluster sum of squared distances to the cluster mean from 619,110 for the convergent log-likelihood identified by hierarchical clustering initialization to 606,377 for the global mode identified by BIA. This example shows that converging to a higher log-likelihood can change the optimal clustering solution of a data set and provide a more intuitive result with clearer separation between groups. Increasingly higher log-likelihoods at convergence may be associated with increasingly more marked divergences in clustering solutions, which at a very minimum should be evaluated for the insights they offer. This demonstrates the importance of converging to the highest possible meaningful log-likelihood and the inherent value of the BIA method.



**Fig. 5** AIS data: histogram of convergent log-likelihoods with  $G = 2$  and ellipsoidal covariance structure with equal shape and volume (EEV) using (a) pyramid burn-in as detailed in (O’Hagan et al, 2012) (b) BIA with 40 initial  $\mathbf{Z}$  matrices and 50 preliminary EM runs (c) BIA with 50 initial  $\mathbf{Z}$  matrices and 100 preliminary EM runs. The purple vertical line at  $-4743.6$  represents the optimal log-likelihood reached by **mclust**’s hierarchical clustering initialization.

#### 4.2 Simulated mixture of Gaussians data


When a multivariate mixture of Gaussians model with the most general covariance structure is fitted to the simulated data as per Baudry et al (2015), hierarchical clustering initialization identifies a 3 group solution as optimal in terms of BIC for 98 of the 100 data sets tested. The same mixture of Gaussian distributions initialized using BIA identifies a 3 group solutions as optimal in terms of BIC in 49 of the 100 cases and a 2 group solution as optimal in the remaining 51 cases. This is preferable in such a case where both the 2 group and 3 group solutions are viable, demonstrating to the user that the underlying horizontal and vertical data structure can be recreated with one less component if so desired. Hence BIA can be seen to offer similar benefits to the Supervised Integrated Complete Likelihood (SICL) decision criterion introduced in Baudry et al (2015), but rather than change the model decision criterion, BIC is retained and instead a different set of decisions is reached via superior initialization of the model.

#### 4.3 Carcinoma data

For each of 100 randomly generated starting values, the EM algorithm was run to convergence for the *Carcinoma* data and multiple local maxima were uncovered. The distribution of convergent lower-bounds to the log-likelihood is shown in Figure 7. Parameter estimates obtained by the LCA model are visualised in Figures 8. These represent the estimated item probabilities  $\hat{\theta}$  as heat maps, with values close to 1 represented by a lighter grey colour, and values close to 0 with darker grey. The estimated class size  $\hat{\tau}$  is represented by the relative length of the cells in each row.

As can be seen in Figure 8, the parameter estimates corresponding to convergent log-likelihood  $-289.79$  and convergent log-likelihood  $-289.29$  are highly similar, however, a different interpretation of the data corresponds to convergent log-likelihood  $-291.27$ . This illustrates the importance of not converging to the sub-optimal mode.






**Fig. 6** AIS data: Pairs plot for body mass index (bmi) and body fat Percentage (Bfat) showing the clustering solution using **mclust**'s hierarchical clustering initialization. The 15 observations that change cluster membership when using the optimal likelihood identified by BIA are overlaid with crosses.

As seen in Table 2, BIA achieves the global maximum convergent likelihood or close to it almost 90% of the time, using 30 random starts and 10 iterations of the algorithm before BIA is applied. This is a marked improvement over the outcome when random starts are employed. Computationally the approach also performed better, with a single run of the BIA algorithm requiring about one third of the time of running to convergence 100 random starting positions.


**Table 2** Distribution of convergent log-likelihoods for 100 initializations of the *Carcinoma* data using BIA and random starts.

Log-likelihood	-294	-293	-292	-291	-290	-289
BIA	0	0	1	10	46	43
Random Starts	1	3	21	12	40	23






**Fig. 7** Histogram of distribution of convergent log-likelihoods using 100 random starts for the *Carcinoma* data.



**Fig. 8** Visualisation of different parameter estimates obtained by local maxima for the LCA model applied to the *Carcinoma* data.



**Fig. 9** Histogram of distribution of convergent log-likelihoods using 100 random starts for the *Alzheimer's* data.

#### 4.4 Alzheimer's data

A 3 group LCA model was applied to the Alzheimer's data from 100 randomly generated starting values. The distribution of the convergent log-likelihoods is shown in Figure 9; the global maximum is only obtained in 18 cases. Table 3 displays the differing solutions arising from applying LCA to the *Alzheimer's* data at convergent log-likelihoods of  $-745.7$  and  $-743.5$  respectively. Again, the latter solution produces a significant alteration in interpretation of symptom composition in the three different classes, with ramifications for subsequent patient classification.

**Table 3** Parameter estimates for the 3 group LCA model fitted to the *Alzheimer's* data. Firstly, the sub-optimal convergent log-likelihood  $-745.7$ ; and secondly, the global convergent log-likelihood  $-743.5$ .

	$\hat{\tau}_g$	Hallucination	Activity	Aggression	Agitation	Diurnal	Affective
1	0.39	0.13	0.70	0.28	0.00	0.30	0.78
2	0.31	0.03	0.46	0.00	0.18	0.04	0.51
3	0.30	0.07	0.80	0.41	1.00	0.38	0.97
1	0.51	0.06	0.52	0.06	0.13	0.10	0.55
2	0.47	0.10	0.79	0.37	0.60	0.37	1.00
3	0.02	0.00	0.82	1.00	0.21	1.00	0.00

The performances of the BIA and annealing methods are compared to that of random starts in Table 4. Each method was run 100 times. In the case of BIA, 20 random starts run for 200 iterations were averaged. In this instance, the approach was computationally slower, a single run of the BIA algorithm run taking about 80% longer than separately running a single run to convergence 100 times.

For the annealing method, after some experimentation, the tuning parameters were set to be  $\nu^0 = 0.12$ ,  $r = 0.87$  and  $s = 10$ . While the BIA approach produces an improved performance, with the global maximum being achieved in 67% of cases, the annealing method converges to the global maximum only once. Despite an extensive search of the tuning parameter space, it proved impossible to substantively improve this performance.

**Table 4** Distribution of convergent log-likelihoods, rounded to the nearest integer, for 100 initialisations of the *Alzheimer's* data using random starts, annealing and BIA.

Log-likelihood	-748	-745	-745	-744	-743
Random Starts	3	51	27	1	18
Annealing	0	97	2	0	1
BIA	0	13	20	0	67

This is demonstrated in Figure 10, which shows heat maps of the convergent log-likelihoods obtained, firstly using deterministic annealing for different settings of  $\nu^0$  and  $r$  and secondly using BIA for different numbers of random starts and EM algorithm iterations per random start. Despite using 400 different parameter combinations, the global mode is rarely discovered using deterministic annealing.<sup>1</sup> Conversely, once the BIA method is allowed to initially run for more than 110 iterations, the global mode is found consistently, suggesting that the clustering performance of BIA is not as sensitive to the choice of initialization parameters.

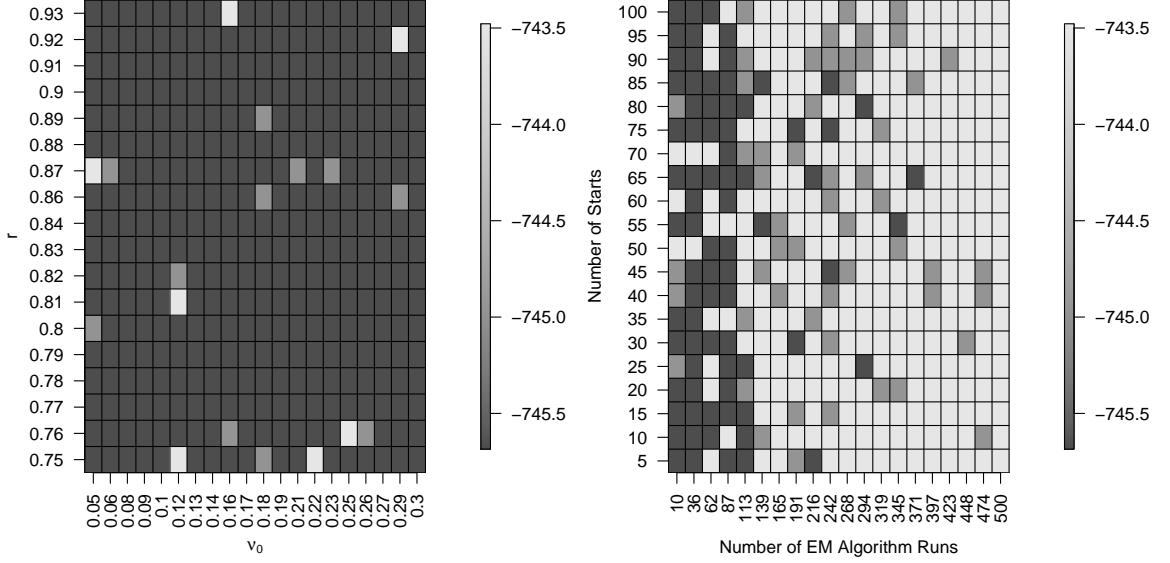
A possible explanation for the poor performance of the annealing method is that in the case of the *Alzheimer's* data the suboptimal solution consists of roughly equal sized groups; the opposite is true in the *Carcinoma* example, where the method's performance is superior to that of BIA. It appears that, unless its tuning parameters are very precisely specified, the annealing method may be biased towards solutions where cluster sizes do not vary significantly. The propensity for the competing methodologies to attain the global maximum convergent log-likelihood is tested across 100 matching random starting positions, with the results presented in Table 4. The BIA approach outperforms the basic random starts and annealing methods.

#### 4.5 Karate data

Obtaining the global maximum for SBM applied to the *Karate* data set when  $G = 4$  using random starts proves to be challenging; over the course of 200 different runs, the maximum was reached only 4 times. A histogram of the various local maxima that can be obtained is shown in Figure 11. Conversely, the BIA approach, using 200 starting values run initially for 15 iterations, obtains the global maximum 19 times out of 20. These results are shown in Table 5. The BIA approach is also much more computationally efficient in this instance; a complete pass of the BIA algorithm is approximately 30 times faster than the process of having a single starting position run to convergence 200 times.

The clustering that obtains the global maximum is visualised in Figure 12. Made clear is the division of the group into two political factions, with each faction having leaders (the nodes labelled 2 and 3, coloured medium-dark and medium-light grey ) and followers (the nodes coloured light

<sup>1</sup> Using  $r = 0.94$  or  $0.95$  regularly resulted in a saddle point with the value of the log-likelihood converging to -770. These results have been omitted from Figure 10.



**Fig. 10** Heat map of convergent log-likelihoods obtained using (left) deterministic annealing with different tuning parameter settings and (right) BIA with different tuning parameter settings, for the *Alzheimer’s* data. A light grey colour indicates an optimal value of the convergent log-likelihood has been obtained.

**Table 5** Distribution of convergent lower bounds to the log-likelihood  $\mathcal{L}$  for the Karate data. For the random starts approach, 200 initialisations were used, while the BIA approach was implemented 20 times.


$\mathcal{L}$	-204	-197	-195	-194	-193	-188	-187	-184	-180
Random Starts	1	34	45	10	83	19	3	1	4
BIA	0	0	0	0	0	0	1	0	19

and and dark grey, labelled 1 and 4), which mainly connect to their respective leaders. Note that despite being the two largest groups, no links are shared between pairs of nodes labelled 1 and 4. This clustering complements the description of the data in Section 2.3.

## 5 Conclusions

We have demonstrated that using BIA across a variety of model-based clustering applications can lead to higher convergent likelihoods being reached and that optimal likelihoods can be reached more frequently than under other competing methods. The importance for the clustering solution of reaching the optimal likelihood was also clearly demonstrated. A sub-optimal convergent log-likelihood can lead to an inferior, potentially misleading clustering solution.

The BIA method improves upon **mclust**’s hierarchical clustering initialization method for the multivariate continuous *AIS* data and reaches a higher convergent likelihood. This higher likelihood in turn changes the clustering solution of the data. While for this example it is not a major




**Fig. 11** Histogram of local maxima obtained for the stochastic blockmodel (SBM) applied to the *Karate* data from 200 random starting values.

change, it does highlight that the clustering solutions obtained by **mclust** are not guaranteed to be optimal. For the simulation experiment using mixtures of overlapping Gaussian components, BIA correctly identifies both relevant solutions whereas hierarchical clustering initialization almost always identifies only the 3 group possibility. For the categorical data examples, the BIA results are reasonably close to those of Zhou and Lange (2010) with respect to the *Carcinoma* data, while the BIA method comprehensively outperforms the annealing approach for the *Alzheimer's* data. For the network data example using the *Karate* data, the BIA approach attains the global maximum likelihood for the 4-group model far more consistently and efficiently than a simple random starts approach.

While the illustrative data sets used were moderate in terms of size and dimensionality, the BIA approach could prove increasingly vital as these quantities increase. This is especially true for network data, where, in its simplest form, the size of a data set increases quadratically with the sample size. Alternatively, the approach may also prove useful for more sophisticated models with additional parameters to be estimated. For example, in the case of continuous data, a mixture of skewed, heavy-tailed distributions such as the skew-t (McLachlan and Peel, 1998; Andrews and McNicholas, 2011; Lee and McLachlan, 2012) may be preferred to a mixture of Gaussians.

An increased number of EM runs and starting positions will generally boost performance of the BIA approach, but as in the case of tempering (Zhou and Lange, 2010) there is no single systematic rule for selecting an optimum combination, and some trial and error is required. However, extensive experimentation on the motivating data sets has shown that the results presented are robust for wide ranges of number of EM runs and number of starting positions. Intuitively it seems reasonable to expect that, if only one of these two inputs is to be elevated, one should favour many starting



**Fig. 12** Plot of the *Karate* network data set, with 4 groups fitted to the data.

positions with a relatively small number of runs for each, since the EM algorithm usually increases likelihood most substantially in the early iterations. For the data sets analyzed, this intuition holds for the *Karate*, *AIS* and *Carcinoma* data, but not the *Alzheimer's* data.

In the examples shown, different starting values were generated using standard approaches. This was done to ensure simple and fair settings were used to compare BIA with competing initialization methods such as random starts, hierarchical clustering, and deterministic annealing. However, this does not have to be the case. For example, in the case of continuous data, clustering solutions obtained from the different covariance structures available in **mclust** could be used. An advantage of the BIA approach is that rather than discarding competing heuristic approaches, they may be assimilated in a principled, statistically robust manner that efficiently harnesses the available computational resources.

**Acknowledgements** The authors would like to acknowledge the contribution of Dr Jason Wyse to this paper, who provided many helpful insights as well as C++ code for the label-switching methodology employed.

## References

- Agresti A (2002) *Categorical Data Analysis*, 2nd edn. John Wiley & Sons
- Aitkin M, Aitkin I (1996) A hybrid EM/Gauss-Newton algorithm for maximum likelihood in mixture distributions. *Statistics and Computing* 6:127–130

- Andrews JL, McNicholas PD (2011) Model-based clustering, classification, and discriminant analysis via mixtures of multivariate t-distributions. *Statistics and Computing* 22(5):1021–1029
- Baudry JP, Celeux G (2015) Em for mixtures. *Statistics and Computing* 25(4):713–726
- Baudry JP, Cardoso M, Celeux G, Amorim MJ, Ferreira AS (2015) Enhancing the selection of a model-based clustering with external categorical variables. *Advances in Data Analysis and Classification* 9(2):177–196, DOI 10.1007/s11634-014-0177-3, URL <http://dx.doi.org/10.1007/s11634-014-0177-3>
- Besag J (1986) On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society Series B* 48:259–302
- Biernacki C, Celeux G, Govaert G (2003) Choosing starting values for the EM algorithm for getting the highest likelihood in multivariate Gaussian mixture models. *Computational Statistics and Data Analysis* 41(3-4):561–575
- Byrd R, Lu P, Nocedal J, Zhu C (1995) A limited memory algorithm for bound constrained optimization. *ACM Transactions on Mathematical Software* 16:1190–1208
- Carpaneto G, Toth P (1980) Algorithm 548: Solution of the assignment problem. *ACM Transactions on Mathematical Software (TOMS)* 6(1):104–111
- Celeux G, Govaert G (1992) A classification em algorithm for clustering and two-stochastic versions. *Computational Statistics and Data Analysis* 14:315–332
- Cook R, Weisberg S (1994) *An Introduction to Regression Graphics*. John Wiley and Sons, New York
- Csardi G, Nepusz T (2006) The igraph software package for complex network research. *InterJournal Complex Systems*:1695
- Daudin JJ, Picard F, Robin S (2008) A mixture model for random graphs. *Statistics and Computing* 18:173–183
- Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B: Methodological* 39(1):1–38, with discussion
- Fraley C, Raftery AE (1999) Mclust: Software for model-based clustering. *Journal of Classification* 16:297–306
- Fruchterman TMJ, Reingold EM (1991) Graph drawing by force-directed placement. *Software - Practice and Experience* 21(11):1129–1164
- Goodman LA (1974) Exploratory latent structure analysis using both identifiable and unidentifiable models. *Biometrika* 61(2):215–231
- Hand DJ, Yu K (2001) Idiot's Bayes: Not so stupid after all? *International Statistical Review / Revue Internationale de Statistique* 69(3):385–398
- Hoeting A, Madigan D, Raftery A, Volinsky C (1999) Bayesian model averaging: a tutorial. *Statistical Science* 14(4):382–401
- Holland PW, Laskey KB, Leinhardt S (1983) Stochastic blockmodels: First steps. *Social Networks* 5(2):109 – 137
- Karlis D, Xekalaki E (2003) Choosing initial values for the EM algorithm for finite mixtures. *Computational Statistics and Data Analysis* 41(3-4):577–590
- Lee S, McLachlan GJ (2012) Finite mixtures of multivariate skew t-distributions: some recent and new results. *Statistics and Computing* 24(2):181–202
- Linzer DA, Lewis JB (2011) poLCA: An **R** package for polytomous variable latent class analysis. *Journal of Statistical Software* 42(10):1–29

- McGrory C, Ahfock D (2014) Transdimensional sequential Monte Carlo for hidden Markov models using variational Bayes - SMCVB. *Proceedings of the 2014 Federated Conference on Computer Science and Information Systems* 3:61–66
- McLachlan GJ, Krishnan T (1997) *The EM Algorithm and Extensions*. Wiley Interscience, New York
- McLachlan GJ, Peel D (1998) *Advances in Pattern Recognition: Joint IAPR International Workshops SSPR'98 and SPR'98 Sydney, Australia, August 11–13, 1998 Proceedings*, Springer Berlin Heidelberg, Berlin, Heidelberg, chap Robust cluster analysis via mixtures of multivariate t-distributions, pp 658–666
- McLachlan GJ, Peel D (2000) *Finite Mixture Models*. Wiley Interscience, New York
- Meng XL, Rubin DB (1992) Recent extensions of the EM algorithm (with discussion). In: *Bayesian Statistics 4*, Oxford University Press, pp 307–320
- Meng XL, Rubin DB (1993) Maximum likelihood estimation via the ECM algorithm: a general framework. *Biometrika* 80:267–278
- Moran M, Walsh C, Lynch A, Coen RF, Coakley D, Lawlor BA (2004) Syndromes of behavioural and psychological symptoms in mild Alzheimer's disease. *International Journal of Geriatric Psychiatry* 19:359–364
- Murphy M, Wang D (2001) Do previous birth interval and mother's education influence infant survival? A Bayesian model averaging analysis of Chinese data. *Population Studies* 55(1):37–47
- Neal RM, Hinton GE (1999) A view of the EM algorithm that justifies incremental, sparse, and other variants. In: Jordan MI (ed) *Learning in graphical models*, MIT Press, Cambridge, MA, USA, pp 355–368
- Nobile A, Fearnside AT (2007) Bayesian finite mixtures with an unknown number of components: The allocation sampler. *Statistics and Computing* 17(2):147–162
- O'Hagan A, Murphy T, Gormley I (2012) Computational aspects of fitting mixture models via the expectation-maximisation algorithm. *Computational Statistics & Data Analysis* 56(12):3843–3864
- Raftery AE, Balabdaoui F, Gneiting T, Polakowski M (2005) Using Bayesian model averaging to calibrate forecast ensembles. *Monthly Weather Review* 133:1155–1174
- Redner R, Walker H (1984) Mixture densities, maximum likelihood, and the EM algorithm. *Society for Industrial and Applied Mathematics Review* 26:195–329
- Rokach L, Maimon O (2010) Clustering methods. In: *Data Mining and Knowledge Discovery Handbook*, Springer, pp 321–352
- Schwarz G (1978) Estimating the dimension of a model. *Annals of Statistics* 6:461–464
- Slonim N, Atwal GS, Tkacik G, Bialek W, Mumford D (2005) Information-based clustering. *Proceedings of the National Academy of Sciences of the United States of America* 102(51):18,297–18,302
- Snijders TAB, Nowicki K (1997) Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *Journal of Classification* 14(1):pp.75–100
- Ueda N (1998) Deterministic annealing EM algorithm. *Neural Networks* 11:271–282
- Volant S, Martin Magniette ML, Robin S (2012) Variational Bayes approach for model aggregation in unsupervised classification with Markovian dependency. *Computational Statistics & Data Analysis* 56(8):2375–2387
- Volinsky CT, Madigan D, Raftery AE, Kronmal RA (1997) Bayesian model averaging in proportional hazard models: assessing the risk of a stroke. *Journal of the Royal Statistical Society Series C (Applied Statistics)* 46(4):433–448



- 
- Walsh C (2006) Latent class analysis identification of syndromes in Alzheimer's disease: A Bayesian approach. *Metodološki Zvezki - Advances in Methodology and Statistics* 3:147 – 162
- Ward JH (1963) Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association* 58(301):236–244
- Wasserman S, Faust K (1994) *Social network analysis: Methods and applications*. Cambridge Univ Press
- White A, Murphy TB (2014) BayesLCA: An R package for Bayesian latent class analysis. *Journal of Statistical Software* 61(13):1–28
- Wintle BA, McCarthy MA, Volinsky CT, Kavanagh RP (2003) The use of Bayesian model averaging to better represent uncertainty in ecological models. *Conservation Biology* 17(6):1579–1590
- Zachary WW (1977) An information flow model for conflict and fission in small groups. *Journal of Anthropological Research* 33(1):452–473
- Zhou H, Lange KL (2010) On the bumpy road to the dominant mode. *Scandinavian Journal of Statistics* 37(4):612–631
- Zhu C, Byrd R, Lu P, Nocedal J (1997) Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *SIAM Journal on Scientific Computing* 23(4):550–560

